

CLAIMS:

1. A method of combinatorial multimodal optimisation for finding multiple optimal ways of dividing a set W of n values into m groups, such that each of
5 the groups satisfies a respective constraint condition, the method comprising:

- (a) defining an initial population of individuals, each representative of a trial solution;
- 10 (b) calculating for each individual a fitness vector indicative of whether the constraint condition for each group has been satisfied;
- (c) selecting a plurality of individuals for the next generation in dependence upon their respective fitness vectors;
- (d) creating a new population including the selected individuals; and
- 15 (e) repeating steps (b) to (d) until the population stabilizes, the individuals of the stable population representing multiple optional ways of dividing the set W .

2. A method as claimed in claim 1 in which the fitness vector is of length m , each element in the fitness vector being indicative of whether the constraint
20 condition of a corresponding one of the m groups has been satisfied.

3. A method as claimed in claim 2 in which the fitness vector comprises m bits, each bit being indicative of whether the constraint condition of a
25 corresponding one of the m groups has been satisfied.

4. A method as claimed in claim 1 including calculating a fitness value for each individual.

5. A method as claimed in claim 4 when dependent upon claim 3 in which the fitness value comprises the sum of the bits in the fitness vector.

5 6. A method as claimed in any one of the preceding claims including reserving a proportion of the new population for individuals selected at step (c).

7. A method as claimed 6 in which a non-reserved proportion of the new population is generated using a Roulette wheel selection method.

10

8. A method as claimed in any one of the preceding claims in which step (c) comprises selecting non-dominated individuals using the criteria of Pareto optimality.

15 9. A method as claimed in claim 8 when dependent upon claim 4 including ranking non-dominated individuals by fitness value, and selecting from the ranked list.

20 10. A method as claimed in claim 9 in which only non-dominated individuals with greatest fitness value may be selected at step (c).

11. A method as claimed in claim 4 in which step (c) comprises selecting individuals in dependence upon both their respective fitness vectors and their respective fitness values.

25

12. A method as claimed in any one of the preceding claims in which crossover and mutation are applied at step (d) to at least some individuals in the new population.

5 13. A method as claimed in any one of the preceding claims in which step (c) comprises selecting no more than one individual for each unique fitness vector.

10 14. A method of distributing a plurality of tasks between a plurality of devices connected together to form a network, wherein each device has an associated constraint on the amount of tasks that it can perform per unit of time, the method comprising:

15 (a) generating a plurality of trial solution allocations to form an initial population of allocations;

(b) calculating for each allocation a fitness vector indicative of whether the constraint condition for each device has been satisfied;

(c) selecting a plurality of allocations for inclusion in the next generation of allocations in dependence upon their respective fitness vectors;

20 (d) creating the next generation of allocations by including the allocations selected in step (c) together with new allocations each of which is formed from a combination of two or more of the allocations selected in step (c);

(e) repeating steps (b) to (d) until the population stabilizes; and

25 (f) allocating the tasks among the devices according to one of the allocations included in the stabilized population.

15. A method as claimed in claim 14 wherein the devices are processors within a multi-processor computer system.

5 16. A method as claimed in either of claims 14 or 15 wherein the devices are computers within a computer network.

17. A method as claimed in claim 14 wherein the devices are routers and the tasks are estimated volumes of traffic to be routed through the routers within a data network, and wherein the allocations are used to form a routing strategy.

10 18. A method as claimed in any one of claims 14 to 17 in which step (c) comprises selecting non-dominated allocations using the criteria of Pareto optimality of the associated fitness vectors.

15 19. A method as claimed in any one of the preceding claims in which new allocations are formed in step (d) by performing crossover operations in respect of groups of two or more of the allocations selected in step (c).

20 20. A method as claimed in any of claims 14 to 19 in which mutation operations are applied to one or more of the new allocations formed in step (d) according to a predetermined probability of each new allocation being mutated.

21. A computer program for carrying out the steps of any one of the preceding claims.

25

22. A carrier medium carrying the computer program of claim 21.

23. A system comprising a plurality of devices connected together to form a network, wherein each device has an associated constraint on the amount of tasks that it can perform per unit of time, the system including means for allocating a plurality of tasks among the devices, the allocation means comprising:

(a) means for generating a plurality of trial solution allocations to form an initial population of allocations;

(b) means for calculating for each allocation a fitness vector indicative of whether the constraint condition for each device has been satisfied;

(c) means for selecting a plurality of allocations for inclusion in the next generation of allocations in dependence upon their respective fitness vectors;

(d) means for creating the next generation of allocations by including the allocations selected in step (c) together with new allocations each of which is formed from a combination of two or more of the allocations selected in step (c);

(e) means for repeating steps (b) to (d) until the population stabilizes; and

(f) means for allocating the tasks among the devices according to one of the allocations included in the stabilized population.